ILC-Americas Workshop

SLAC October 14-16, 2004

LLRF/Global Phase Control

Larry Doolittle (LBNL) and Stefan Simrock (DESY)

Outline

- What is the scope of LLRF/Global Phase Control?
- Survey of existing technology
- Assessment of challenges facing project
- Conclusions

Scope of LLRF/Global Phase Control

- ∼600 klystron control systems, each with
  - 100 precision 1300 MHz RF vector receivers
  - 32 high voltage piezoelectric drivers
  - 96 waveguide tuner DC motor controllers
  - 32 cavity tuner stepper motor controllers
  - 1 klystron drive signal output
- Programming to drive it
  - General purpose computer
  - Digital Signal Processor (DSP)
  - Field Programmable Gate Array (FPGA)
- Software tasks:
  - Feedback within pulse: klystron output and piezo tuners
  - pulse-to-pulse feedback
  - slow tuning of cavities and waveguides
  - self-calibration and drift correction

Requirements for RF Control

- Derived from beam properties
  - energy spread
  - emittance
  - bunch length (bunch compressor)
  - arrival time
- Different accelerators have different requirements for field stability (approximate RMS requirements)
  - 1% for amplitude and 1° for phase (example: SNS)
  - 0.1% for amplitude and 0.1° for phase (linear collider)
  - up to 0.01% for amplitude and 0.01° for phase (XFEL)

Note: Distinguish between correlated and uncorrelated error

(from Stefan Simrock, Linac 2004)

Sources of Perturbations

o Beam loading - Beam current fluctuations - Pulsed beam transients - Multipacting and field emission - Excitation of HOMs - Excitation of other passband modes - Wake fields o Cavity drive signal - HV- Pulse flatness - HV PS ripple - Phase noise from master oscillator - Timing signal jitter - Mismatch in power distribution	o Cavity dynamics - cavity filling - settling time of field o Cavity resonance frequency change - thermal effects (power dependent) - Microphonics - Lorentz force detuning o Other - Response of feedback system - Interlock trips - Thermal drifts (electronics, power amplifiers, cables, power transmission system)

System integration

- high power RF (klystron nonlinearity, cathode voltage regulation)
- phase and frequency reference infrastructure
- cavity hardware (mechanical resonances, piezoelectric actuator)
- global control system (hardware, software interfaces)
- tunnel design (cable/waveguide latency, radiation environment)
- power and cooling

Survey of existing technology

Not hard to predict what RF/electronics technology will be available in a few years, and in fact what we have in-hand is adequate to do the job. Some of the expected technology enhancements will help, but won't change the basic requirements or techniques.

Digital Control at the TTF

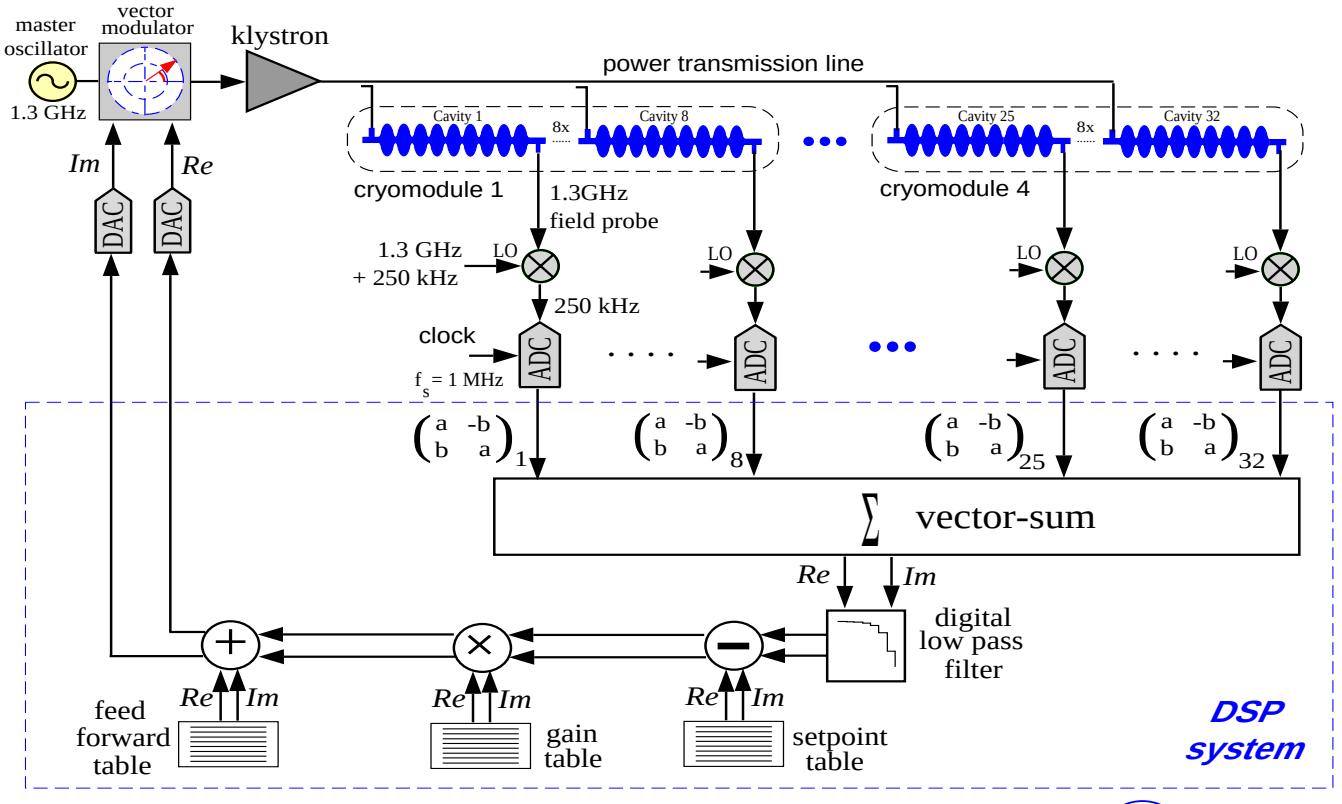

DSP and ADC board LLRF for TTF I

Performance at TTF (1)

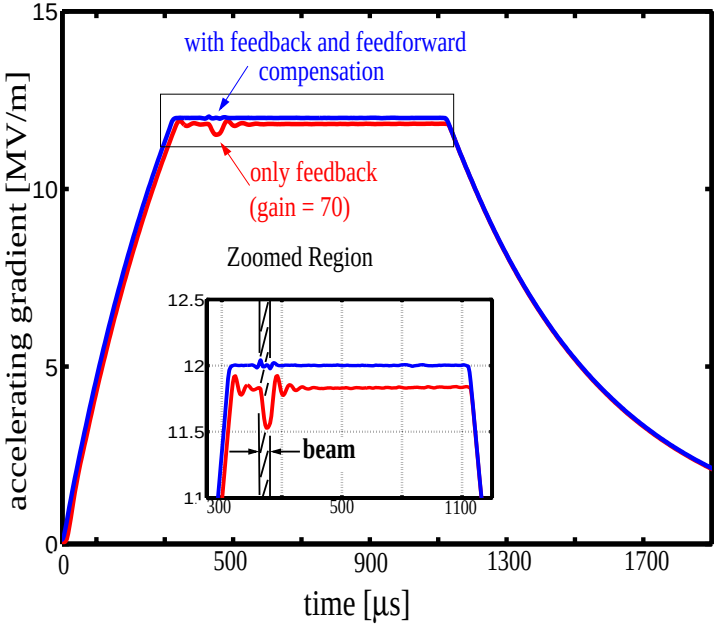

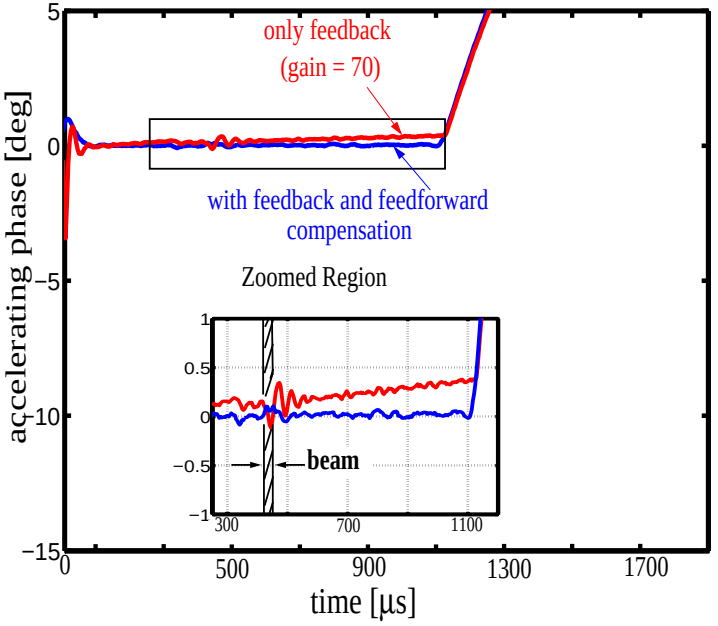

Amplitude **Phase**

New Field Control Module for SNS

The prototype Field Control Module. The Analog Front End, Digital Front End, and RF Output daughtercards (left to right) are mounted on the VXI motherboard .

L. Doolittle

Matthias Liepe
CORNELL
UNIVERSITY

Ultra-Fast Digital RF Field Control System for CESR and ERLs

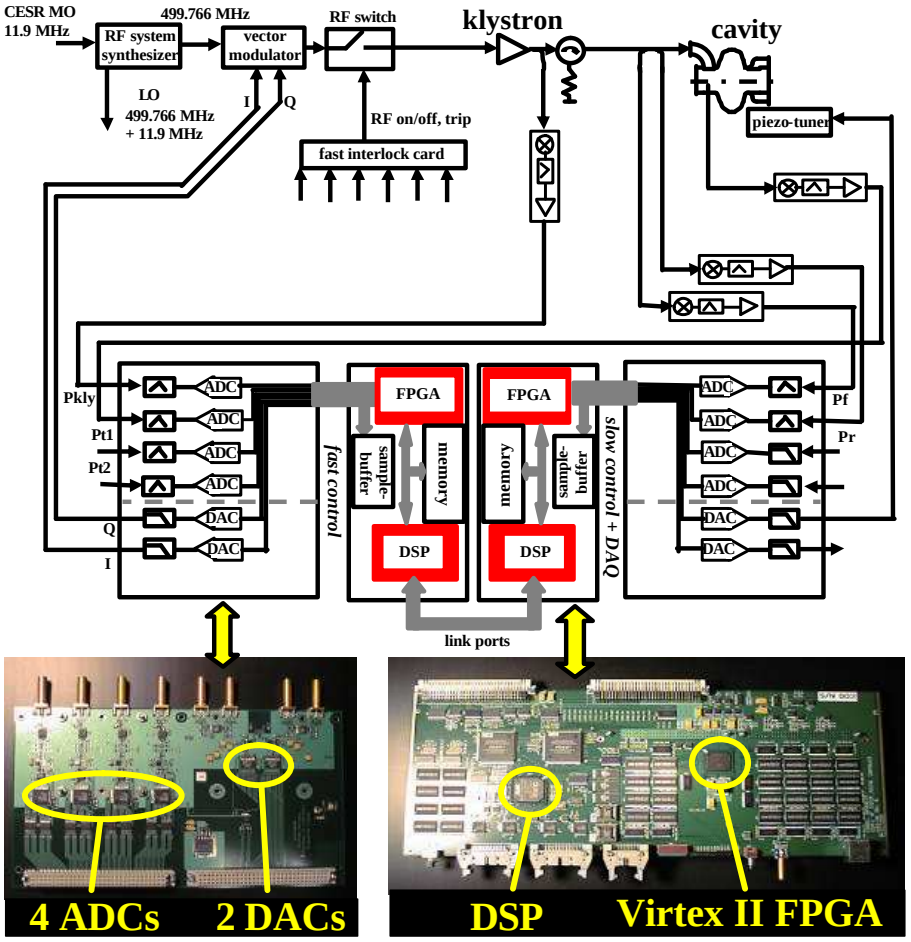

- *very low delay in the control loop (≈ 1 μs)*
- *Field Programmable Gate Array (FPGA) design combines the speed of an analog system and the flexibility of a digital system*
- *high computation power allows advanced control algorithms*
- *all boards have been designed in house*
- *generic design: digital boards can be used for a variety of control and data processing applications*

Active Compensation of Lorentz Force Detuning (2)

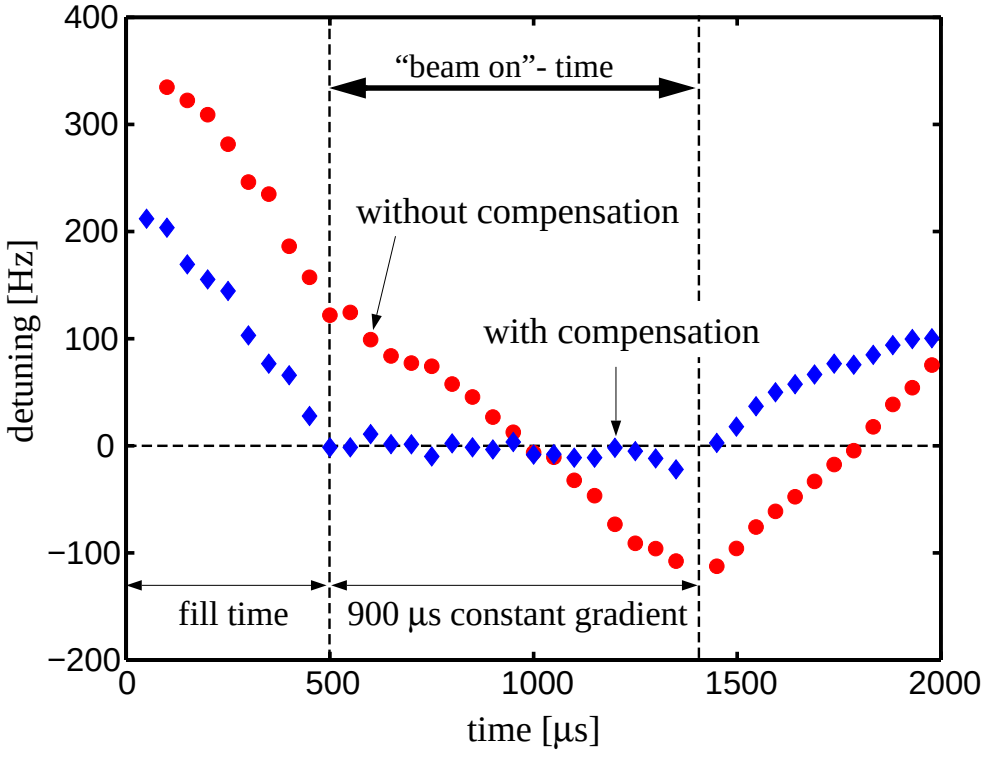

9-cell cavity operated at 23.5 MV/m

Lorentz force compensated with fast piezoelectric tuner

Comparison of available fast ADCs

part	ADS5500	ADS809	AD6645	AD9430	AD9433	AD9236	CLC5957	MAX1418	MAX12553	
density	1	1	1	1	1	1	1	1	1	channels
resolution	14	12	14	12	12	12	12	15	14	bits
speed	125	80	105	210	125	80	70	65	65	MS/s
input 3dB BW	750	1000	270	700	750	500	300	260	400	MHz
aperture jitter	0.30	0.25	0.10	0.25	0.25	0.30	0.30	0.21	0.2	ps rms
latency	16.5	5	3.5	14	10	7	3	3	8.5	cycles
min. latency	132.0	62.5	33.3	66.7	80.0	87.5	42.9	46.2	130.8	ns
power	780	900	1500	1300	1350	366	640	2000	363	mW
power (idle)	180	20	N/A	N/A	N/A	1.0	N/A	N/A	0.02	mW
gain tempco	100	20	48	200	-125	12	?	?	?	ppm/°C
analog V_{CC}	3.3	5.0	5.0	3.3	5.0	3.3	5.0	5.0	3.3	V
digital V_{CC}	3.3	3.3	3.3	3.3	3.3	3.3	5.0	2.5-3	2.0	V
pins	64	48	52	100	52	32	48	56	40	
size	12x12	9x9	12x12	16x16	12x12	5x5	12x8	8x8	6x6	mm
cost	123.50	32.44	88.00	107.06	76.00	29.74	19.10	24.85	24.85	US$
clock	diff	diff	diff	diff	diff	SE	diff	diff	diff	

prices as of 2004-04-22, generally in quantity 100

- Latency is not going down with modern parts!
- Hope for 14-bit, 2-channel, 80 MS/s parts with 300 mW dissipation per converter

Assessment of challenges facing project

• Basic understanding of physics ("plant") and electronics and control system theory is not a problem.

• System integration, availability engineering (especially with a single tunnel, wow!), involves real work.

• Managing complexity, and enforcing simple solutions to simple problems, is both essential and incredibly difficult as the scale of the installation reaches the level needed for ILC.

  ◦ subdivide the problem with rational and well understood interfaces

RFC-1925 - The Fundamental Truths

April 1, 1996

ard you push and no matter what the priority, you can't increase the speed

. No matter how hard you try, you can't make a baby in much less than 9
rying to speed this up *might* make it slower, but it won't make it happen
er.

rust, pigs fly just fine. However, this is not necessarily a good idea. It is
where they are going to land, and it could be dangerous sitting under them
ead.

fe can never be fully appreciated nor understood unless experienced
things in ~~networking~~ ACCELERATORS can never be fully understood by
either builds commercial ~~networking~~ ACCELERATOR equipment nor runs
~~etwork~~ ACCELERATOR.

ible to agglutinate multiple separate problems into a single complex
olution. In most cases this is a bad idea.

RFC-1925 - The Fundamental Truths – cont.

ve a problem around (for example, by moving the problem to a different
ll ~~networking~~ ACCELERATOR CONTROLS architecture) than it is to solve

. It is always possible to add another level of indirection.

ething
. Good, Fast, Cheap: Pick any two (you can't have all three).

cated than you think.

, whatever it is, you need more.
Every ~~networking~~ ACCELERATOR problem always takes longer to solve
ems like it should.

ver fits all.

ea will be proposed again with a different name and a different
ardless of whether it works.
. See rule 6a.

LRF design, perfection has been reached not when there is nothing left to
ere is nothing left to take away.

Familiar block diagram

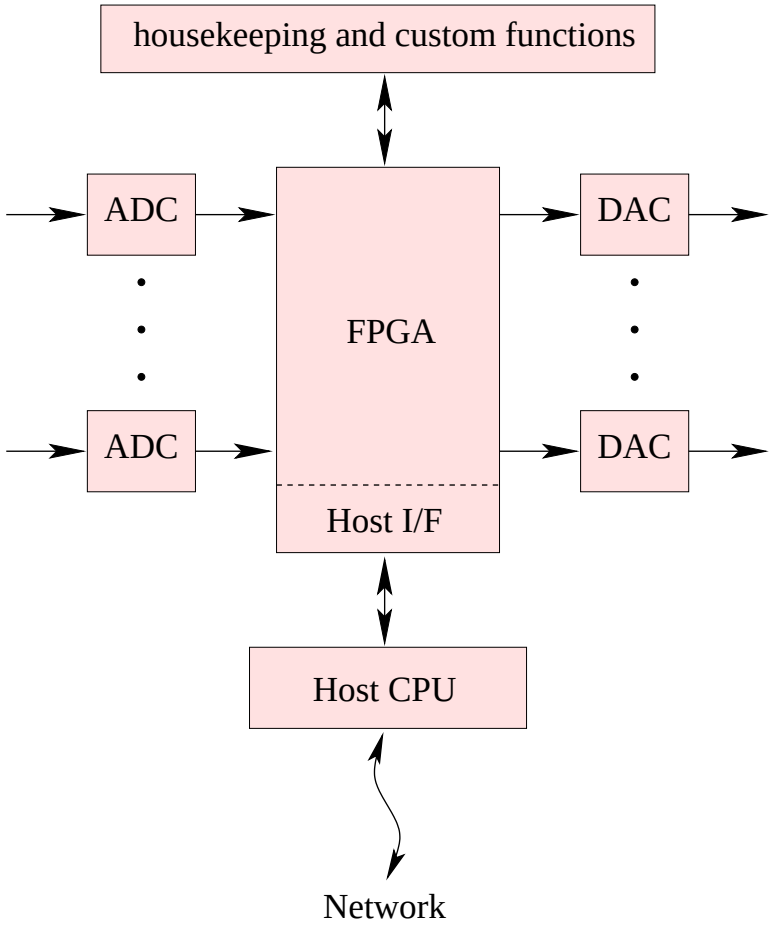

Familiar block diagram, scaled for ILC LLRF

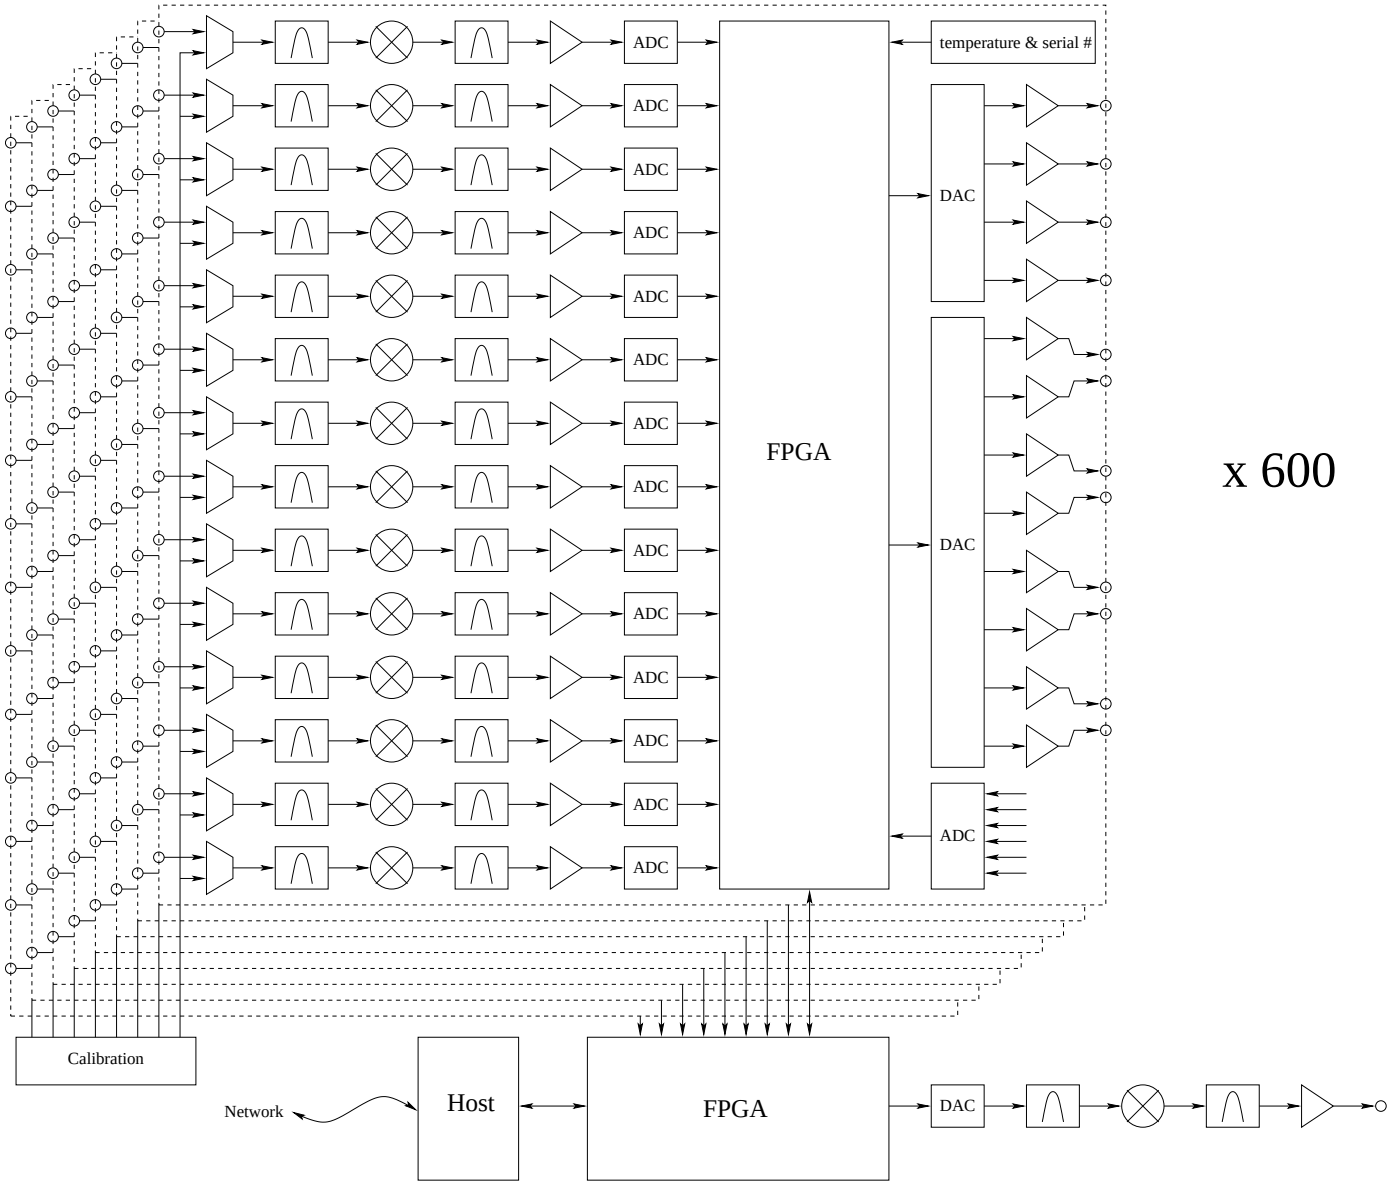

x 600

Commentary

• The road from concept to reality has so many traps that it would be foolish to claim it's "just engineering", even though it is.

• Software: huge amount of engineering gets sucked up into this category.

• Old joke: What's the difference between hardware and software?

  ◦ Hardware keeps getting cheaper, faster, and smaller!

• Scale models, simulation stands, verification/testing techniques for hardware and software. Put a cavity model in the production FPGA design, lets you operate the whole linac before the klystrons are turned on. Make sure all the upper level software works under full load.

• Latency (group delay) nominal budget:

300 ns	klystron
300 ns	cables
100 ns	filters
300 ns	ADC+FPGA+DAC
1000 ns	total

Technical Questions

When the beam sits in the damping ring (for 200 ms), its charge is measured; that info can be passed to the linac, and get to the controller before the bunch hits it. Maybe get that information 100 ms in advance? How?

First time (after each shutdown) turn-on is difficult, since nobody thinks you can dead-reckon the phase, and a stray bunch can burn a hole in the vacuum pipe. Can artificially degrade emittance of first bunch so it can be used to determine initial phasing. Some ideas to propagate a 5 GeV pilot bunch with cavities turned off.

Trade off location of mixer and ADC relative to distance from cavity - calibration difficulty, latency, radiation hardness.

Use air conditioning or not? Need a strategy for calibration, temperature monitoring. Consider capital, installation, reliability, and operating costs of cooling.

What fraction of development, design, infrastructure, and hardware can be shared with the BPM system?

Technical Questions (continued)

Is there a commercial crate standard (e.g., VXI, cPCI) that is appropriate for application at this scale? Personal prejudice: NO; use gigalink or other network cables, and hot-swappable self-contained modules. It's quite a trick to pick standards that will be supported by industry for the 20+ year lifetime of an accelerator. Hope to minimize barriers to upgrading components.

What are this accelerator's requirements for a global control system? Does existing software meet these needs? At what level does a software-intensive subsystem like LLRF connect with it? How is this interface designed and tested?

How close are we expected to run to klystron saturation? Energy efficiency, power supply cost, tradeoff *vs.* control system robustness. The closer you try to run a system to its limits, the more complex it is to handle.

There would be a real advantage to spitting out a test pulse or two in the 200 ms between pulses, to measure the mechanical state of the cavity. Can the capacitor bank tolerate occasional 10 μs pulses in the time leading up to the main driving pulse? Backup plan is to use a piezo vibration sensor, may or may not be well correlated with cavity detuning.

Technical Questions (continued)

Define requirements and suggested interfaces to cover accelerator-wide hard-real-time communications, including interlocks, timing, and mode selection. In a properly designed accelerator timing/communications system, there are no "surprise" asynchronous events, all ∼6000 FPGAs run synchronously.

Is there sufficient benefit to collaboration between adjacent modules within a pulse – e.g., load sharing in case of soft quench – to justify the additional complexity, especially in algorithms (hardware isn't that difficult).

Radiation tolerance - If the LLRF is part of a "single tunnel" design, tremendous extra resources will need to be expended to understand how to operate in a moderate radiation field. Single-event-upset (SEU) tolerance, total ionizing dose component lifetime limit (on-board dosimetry to allow scheduling module exchange).

Redundancy - need operation at reduced specs when an input channel breaks. Single points of failure need special attention: output channel.

Which group in ILC will lead robotic application development? Maintaining a machine at this scale without robots would be a nightmare.

Ideal maintenance cycle

1. In-situ diagnostics to determine which module has degraded/failed
2. Information presented to operators in control room
3. Workarounds automatically engaged to keep accelerator running as well as possible in the face of all but the worst failure modes
4. Robotic exchange of module in tunnel during daily maintenance minute
5. New module is tested in-situ, presumably by firing an RF pulse without beam
6. New module works without operator intervention, using calibration information combined from new module bench measurements, old module plant characterization tables, and one calibration beam pulse

Summary and Conclusions

- The community has a good understanding of the appropriate control techniques for a klystron/cavity system of this type
- Available electronics components are adequate to implement these techniques
  - assumes electronics are not subjected to high radiation
- Software and system integration needs are immense
  - many traps lurk for the unwary
- No technical problems, only difficulty is finding the right people. Of course, if you have the wrong people, you find technical problems.